

DANCo: Dimensionality from Angle and Norm Concentration

C. Ceruti, S. Bassis, A. Rozza, G. Lombardi, E. Casiraghi, and P. Campadelli

Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano,
Via Comelico 39-41, 20135 Milano, Italy
claudio.ceruti@unimi.it

Abstract. In the last decades the estimation of the intrinsic dimensionality of a dataset has gained considerable importance. Despite the great deal of research work devoted to this task, most of the proposed solutions prove to be unreliable when the intrinsic dimensionality of the input dataset is high and the manifold where the points lie is nonlinearly embedded in a higher dimensional space. In this paper we propose a novel robust intrinsic dimensionality estimator that exploits the twofold complementary information conveyed both by the normalized nearest neighbor distances and by the angles computed on couples of neighboring points, providing also closed-forms for the Kullback-Leibler divergences of the respective distributions. Experiments performed on both synthetic and real datasets highlight the robustness and the effectiveness of the proposed algorithm when compared to state of the art methodologies.

Keywords: Intrinsic dimensionality estimation, manifold learning, von Mises distribution, Kullback-Leibler divergence.

1 Introduction

Given a dataset $\mathbf{X}_N \equiv \{\mathbf{x}_i\}_{i=1}^N \subset \mathbb{R}^D$, its intrinsic dimension (**id**) is the minimum number of parameters needed to represent the data without information loss. In the last decade a great deal of research work has been devoted to the development of **id** estimation algorithms; to this aim, the feature vectors \mathbf{x}_i are generally viewed as points constrained to lie on a low dimensional manifold $\mathcal{M} \subseteq \mathbb{R}^d$ embedded in a higher dimensional space \mathbb{R}^D , where d is the **id** to be estimated. In more general terms, according to [16], \mathbf{X}_N is said to have **id** equal to $d \in \{1..D\}$ if its elements lie entirely within a d -dimensional subspace of \mathbb{R}^D .

The **id** is a very useful information for the following reasons. At first, dimensionality reduction techniques, which are often used to reduce the “curse of dimensionality” effect [21] by computing a more compact representation of the data, are profitable when the number of projection dimensions is the minimal one that allows to retain the maximum amount of useful information expressed by the data. Furthermore, when using an auto-associative neural network [23] to perform a nonlinear feature extraction, the **id** can suggest a reasonable value for

the number of hidden neurons. Moreover, according to the statistical learning theory [38], the capacity and the generalization capability of a classifier may depend on the id . In particular, in [14] the authors mark that, in order to balance a classifier’s generalization ability and its empirical error, the complexity of the classification model should also be related to the id of the available dataset. Finally, as it has been recently shown in [4], id estimation methods are used to evaluate the model order in a time series, which is crucial to make reliable time series predictions; this consideration is supported by the fact that the domain of attraction of a nonlinear dynamic system has a very complex geometric structure and the studies on the geometry of the attraction domain are closely related to fractal geometry, and therefore to fractal dimension.

Unfortunately, even if a great deal of research work has been focused at the development of id estimators, and several interesting techniques have been presented in the literature, to our knowledge only few methods [5,28,34,33] have investigated the problem of input datasets having a sufficiently high id (that is $\text{id} \geq 10$) and being drawn from manifolds nonlinearly embedded in higher dimensional spaces; this fact is also highlighted by the experiments reported in this paper showing that well-known techniques fail when dealing with this kind of data. More precisely, it can be noted that several methods underestimate the id if its value is too high. These considerations lead us to the development of an id estimator, called “DANCo” (Dimensionality from Angle and Norm Concentration), that is less affected by underestimation problems, as it is shown by experiments on both synthetic and real datasets, and by the comparison of the achieved results with those reported by state of the art algorithms. The peculiarities and strengths of the proposed estimator are to be sought in the joint use of normalized nearest neighbor distances and mutual angles, whose coupled exploitation allows to reduce the effects of well-known problems such as curse of dimensionality, edge effect, and overall orthogonality.

This paper is organized as follows: in Section 2 previous works on id estimators are reviewed. In Section 3 base theoretical results laying foundations for the proposed estimator are presented. Section 4 sketches the proposed algorithm providing a concise analysis of its properties. A detailed comparison with state of the art methodologies on a wide family of datasets is summarized in Section 5. Finally, Section 6 reports conclusions and future works.

2 Related Works

In this section we summarize the literature related to id estimation methods; note that a more detailed description is reported in the survey [3].

The most cited id estimator is the Principal Component Analysis (PCA) [22], which projects the input dataset on the d directions of maximum variance (principal components, PCs). Exploiting PCA, the estimated id is the number of PCs whose corresponding normalized eigenvalues are higher than a thresholding parameter, usually difficult to be set. More accurate results can be obtained by applying a local PCA [15] that determines the id by combining local estimates

computed in small subregions of the dataset; unfortunately, complications arise in the identification of local regions and in the selection of thresholds [39]. In [1] Bishop describes a Bayesian treatment of PCA (BPCA) to automatically estimate the id of the input dataset. This technique has been extended in [27] to cope with exponential family distributions, but this method requires the knowledge of the distribution underlying the data. To achieve an automatic selection of meaningful PCs, in [18] the authors propose the Sparse Probabilistic Principal Component Analysis (SPPCA) that exploits the sparsity of the projection matrix through a probabilistic Bayesian formulation. PCA-based methods, such as those previously mentioned, are usually classified as projection methods [3,26] since they search for the best subspace where to project the data; unfortunately, they cannot provide reliable id estimates since they are too sensitive to noise and parameter settings [26].

Geometric id estimation methods [26] are most often based on some statistics related to either the distances between neighboring points or the fractal dimension, expressing them as functions of the id of the embedded manifold. The most popular fractal dimension estimator is the Correlation Dimension (CD) [17] that is based on the assumption that the volume of a d -dimensional set scales as r^d with its size r . Since the performances of CD are affected by the choice of the scale r , in [19] the author suggests an estimator (here called **Hein**) based on the asymptotes of a smoothed version of the CD estimate. In [11] the authors present an algorithm to estimate the id of a manifold in a small neighborhood of a selected point, and they analyze its finite-sample convergence properties. Another technique, based on the analysis of point neighborhoods, is the Maximum Likelihood Estimator (MLE) [26] that applies the principle of maximum likelihood to the distances between neighboring points. In [8] the authors propose an algorithm that exploits entropic graphs to estimate both the id and the intrinsic entropy of a manifold; they test their method by adopting either the Geodesic Minimal Spanning Tree (GMST [7]), where the arc weights are the geodesic distances computed through the ISOMAP algorithm [36], or the more efficient kNN-graph (kNNG [8]), where the arc weights are based on Euclidean distances.

We note that many neighborhood based estimators usually underestimate the id when its value is sufficiently high and, to our knowledge, only few works address this problem [5,34,28]. Indeed, as shown in [10], the number of sample points required to perform dimensionality estimation grows exponentially with the id (“curse of dimensionality”). For this reason, when the dimensionality is too high the number of sample points practically available is insufficient to compute an acceptable id estimation. Moreover, the ratio between the points close to the edge of the manifold and the points inside it raises in probability when the dimensionality increases (“edge effect”, [39]), affecting the results achieved by estimators based on statistics related to the behavior of point neighborhoods.

In [5], the authors propose an empirical id correction procedure based on the estimation of the error obtained on synthetically produced datasets of known dimensionality. More precisely, after generating D datasets characterized by in-

cremental **id** values ($d_i \in \{1..D\}$), the authors apply the **CD** algorithm [17] to estimate the **id** (\hat{d}_i) of each dataset. Fitting the points (d_i, \hat{d}_i) they obtain the so-called “correction curve” used to adjust the **id** estimates. In [34] a local estimator (called **IDEA**) based on an asymptotic correction is proposed. To this aim, given a dataset of unknown **id**, random subsets of different cardinalities are extracted and their **id** estimates are computed; the bi-dimensional points composed by the cardinality of each subset and by its **id** estimate are then fitted with a curve having a horizontal asymptote whose ordinate is the final **id**. In [28] the authors describe a method (called **MIND_{KL}**) based on the comparison between the empirical probability density function of the neighborhood distances computed on the dataset and the distribution of the neighborhood distances computed from points uniformly drawn from hyperspheres of known increasing dimensionality; the **id** estimate is the one minimizing the Kullback-Leibler divergence (KL).

3 Theoretical Results

Consider a manifold $\mathcal{M} \equiv \mathbb{R}^d$ embedded in a higher dimensional space \mathbb{R}^D through a locally isometric nonlinear smooth map $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^D$; to estimate the **id** of \mathcal{M} by means of points drawn from the embedded manifold through a smooth probability density function (**pdf**) f , we need to identify a “mathematical object” depending only on d , and we should define a consistent estimator for d based on it.

Assume by hypothesis that the employed manifold sampling process is driven by a smooth **pdf** f ; moreover, consider a spherical neighborhood of the origin $\mathbf{0}_d$ having radius ϵ ; denoting with $\chi_{\mathcal{B}_d(\mathbf{0}_d, 1)}$ the indicator function on the unit ball $\mathcal{B}_d(\mathbf{0}_d, 1)$, the **pdf** restricted to such a neighborhood is:

$$f_\epsilon(\mathbf{z}) = \frac{f(\epsilon\mathbf{z})\chi_{\mathcal{B}_d(\mathbf{0}_d, 1)}(\mathbf{z})}{\int_{\mathbf{t} \in \mathcal{B}_d(\mathbf{0}_d, 1)} f(\epsilon\mathbf{t})d\mathbf{t}} \quad (1)$$

In [28] the authors prove the following:

Theorem 1. *Given $\{\epsilon_i\} \rightarrow 0^+$, Equation (1) describes a sequence of **pdfs** having the unit d -dimensional ball as support; such sequence converges uniformly to the uniform distribution \mathbf{B}_d in the ball $\mathcal{B}_d(\mathbf{0}_d, 1)$.*

Theorem 1 ensures that, from a theoretical standpoint, in our setting it is possible to assume uniformly distributed points in every neighborhood of \mathcal{M} ; in other words, it is possible to define consistent estimators based on local information, assuming without loss of generality that the normalized points are uniformly drawn from the unit hypersphere.

Our technique exploits the statistical properties of norms and mutual angles computed on points drawn from uniformly sampled hyperspheres; to this aim, in Sections 3.1, 3.2 we sketch the statistical properties of norms and angles respectively, while in Section 3.3 we describe how both the above properties can be simultaneously used to define a consistent estimator of the manifold’s **id**.

3.1 Concentration of Norms

Consider initially the problem of estimating the **id** of \mathcal{M} by means of a sample $\{\mathbf{z}_i\}_{i=1}^k$ of points uniformly drawn from $\mathcal{B}_d(\mathbf{0}_d, 1)$; to this aim, we exploit the concentration of norms that is dimensionality-dependent.

In [28] it is shown that the **pdf** associated to the normalized distance r between the hypersphere center and its nearest neighbor is the following:

$$g(r; k, d) = kdr^{d-1}(1 - r^d)^{k-1} \quad (2)$$

Theorem 1 proves that the convergence of f_ϵ to \mathbf{B}_d is uniform, so that when $\epsilon \rightarrow 0^+$ the **pdf** related to the geodetic distances $\frac{1}{\epsilon}\delta_\phi(\phi(\mathbf{0}_d), \phi(\mathbf{z}))$ converges to the **pdf** g defined in Equation (2). Notice that, once k is fixed, $\mathcal{G} = \{g(r; k, d)\}_{d=1}^D$ represents a finite family of D **pdfs** for all the parameter values $1 \leq d \leq D$.

As reported in [28], a Maximum Likelihood estimator (ML) could be found for the parameter d of g , but the resulting estimate may be poor due to the usage of the **kNN** algorithm. More precisely, in high dimensional spaces, the **kNN** method is strongly affected by the edge effect [39] that reduces the quality of the neighborhood estimation.

To obtain a more reliable estimate of d , in [28] the authors propose to minimize the KL divergence between the **pdf** computed on the dataset and those calculated on synthetic data of known **ids**; to this aim, they adopted the KL method proposed in [41].

However, though this KL approach can be applied to every dataset without any restriction on the underlying distribution, in our problem the closed-form for the KL divergence between two minimum neighbor distance **pdfs** can be analytically identified. To this aim, once the parameter k is fixed, we need to estimate the parameter d in g ; to accomplish this task, we decided to employ the ML estimator proposed in [28]. Calling \hat{d}_{ML} the ML estimation obtained on the dataset, and $\check{d}_{d,ML}$ the ML estimations obtained by means of points sampled from d -dimensional hyperspheres¹ (for $d \in \{1..D\}$), the closed-form of the KL for the minimum neighbor distances is:

$$\begin{aligned} \overline{KL}_d &= \mathcal{KL}(g(\cdot; k, \hat{d}_{ML}), g(\cdot; k, \check{d}_{d,ML})) = \int_0^1 g(r; k, \hat{d}_{ML}) \log \left(\frac{g(r; k, \hat{d}_{ML})}{g(r; k, \check{d}_{d,ML})} \right) dr \\ &= \mathcal{H}_k \frac{\check{d}_{d,ML}}{\hat{d}_{ML}} - 1 - \mathcal{H}_{k-1} - \log \frac{\check{d}_{d,ML}}{\hat{d}_{ML}} - (k-1) \sum_{i=0}^k (-)^i \binom{k}{i} \Psi \left(1 + \frac{i\hat{d}_{ML}}{\check{d}_{d,ML}} \right) \end{aligned} \quad (3)$$

where $\mathcal{KL}(\cdot, \cdot)$ is the KL divergence operator, \mathcal{H}_k represents the k -th harmonic number $\left(\mathcal{H}_k = \sum_{i=1}^k \frac{1}{i}\right)$, and $\Psi(\cdot)$ is the digamma function.

¹ Notice that, due to the **kNN** bias effect described above, the ML estimates $\check{d}_{d,ML}$ are biased w.r.t. the real value d employed in the sampling process, and a similar bias can be observed also in the estimated \hat{d}_{ML} .

3.2 Concentration of Angles

As it happens for norms, in high dimensions pairwise angles among k uniformly distributed unitary vectors $\{\mathbf{x}_i\}_{i=1}^k$ on a $(d-1)$ -dimensional surface S^{d-1} of a hypersphere in \mathbb{R}^d are subject to the concentration of their values. The common belief that in high dimensions such vectors tend to be orthogonal to each other has found partly theoretical justification in the past [30], but only in the last decades an even deeper investigation has allowed a more precise characterization of this fact [35].

Two of the most common distributions adopted in circular and directional statistics are the von Mises distribution (VM) and its high-dimensional generalization termed von Mises-Fisher distribution (VMF). More precisely, for $\mathbf{x} \in S^{d-1}$, the VMF distribution has the following form:

$$q(\mathbf{x}; \boldsymbol{\nu}, \tau, d) = C_d(\tau) \exp(\tau \boldsymbol{\nu}^T \mathbf{x}) \quad (4)$$

where the unit vector $\boldsymbol{\nu}$ denotes the mean direction, and the concentration parameter $\tau \geq 0$ gets high values in case of a high concentration of the distribution around the mean direction. In particular, $\tau = 0$ when points are uniformly distributed on S^{d-1} . Moreover, the normalization constant $C_d(\tau)$ in Equation (4) takes the following form:

$$C_d(\tau) = \frac{\tau^{d/2-1}}{(2\pi)^{d/2} I_{d/2-1}(\tau)}$$

where I_v is the modified Bessel function of the first kind with order v . Due to the normalization factor, this pdf is difficult to be used in theoretical derivations; moreover, in the assumptions of Theorem 1, no information about d may be estimated by the knowledge of parameters $(\boldsymbol{\nu}, \tau)$, being $\boldsymbol{\nu}$ uninformative when the hypersolid angles are uniformly distributed ($\tau = 0$), which is the case of a uniformly sampled hypersphere.

Therefore, to infer the id of \mathcal{M} by exploiting the angular information, we focused on the distribution of the angles θ computed between independent pairs of random points chosen in neighborhoods of \mathbb{R}^d and sampled from the uniform distribution in the hypersphere. Note that working on pairwise angles allows both to exploit the concentration factor τ , which is strictly related to the dimensionality d as we will show, and to rely on the VM distribution, which is more tractable w.r.t. the VMF pdf.

With the above notation, considering the angle $\theta \in [-\pi, \pi]$ between two vectors, the VM pdf of θ reads as:

$$q(\theta; \nu, \tau) = \frac{e^{\tau \cos(\theta - \nu)}}{2\pi I_0(\tau)} \chi_{[-\pi, \pi]}(\theta) \quad (5)$$

with the same parameters and notation adopted for the VMF pdf. Intuitively, the VM distribution is the circular counterpart of the normal distribution on a line, sharing with the latter many interesting properties [2]. To understand the

link between τ and d , we firstly recall that $q(\theta; \nu, \tau)$ is unimodal for $\tau > 0$, as a Gaussian random variable peaked around its mean. Next, according to the following theorem, increasing values of τ are expected for points uniformly drawn from hyperspheres with increasing dimensionality d .

Theorem 2. *Given two independent random unit vectors $(\mathbf{x}_1, \mathbf{x}_2)$ in \mathbb{R}^d , chosen from a uniform distribution on S^{d-1} , the concentration parameter τ of the VM distribution describing the angle θ between \mathbf{x}_1 and \mathbf{x}_2 converges asymptotically to the dimensionality d .*

Proof. Consider the following results:

- i) for large concentration values τ , a VM distribution with parameters (ν, τ) becomes a Gaussian distribution with mean ν and standard deviation $1/\sqrt{\tau}$ [20];
- ii) performing the variable substitution $\theta = \sqrt{d}(\theta - \pi/2)$, the resulting random variable converges in distribution to a standard normal one [35].

Combining i) and ii), it follows that θ asymptotically follows a Gaussian pdf with mean $\nu = \pi/2$ and standard deviation $\sigma = 1/\sqrt{\tau} = 1/\sqrt{d}$, which holds only when $\tau = d$.

Theorem 2 has both a general and a specific value. At first, it formally proves the existence of the concentration of angles in high dimensions, stating both an asymptotic linear relation between concentration and dimensionality, and the orthogonality between any couple of infinite-dimensional vectors. Moreover, Theorem 2 allows to estimate the id (d) of the observed points through the estimation of the concentration parameter τ .

The methodology we propose in Section 4 employs both the ML estimation of the VM parameters ν and τ , and the KL divergence between the VM pdf estimated from the observed dataset and those computed on synthetic data of known ids. Assuming that $\{\theta_1, \dots, \theta_N\}$ is a sample drawn from a VM distribution with parameters (ν, τ) , the ML of the population direction ν equals the sample mean direction; more precisely:

$$\hat{\nu} = \arctan \frac{\sum_{i=1}^N \sin \theta_i}{\sum_{i=1}^N \cos \theta_i} \quad (6)$$

Likewise, the ML of the concentration parameter τ equals the concentration parameter $\hat{\tau}$ calculated as a solution of $\eta = \frac{I_1(\tau)}{I_0(\tau)} \equiv A(\tau)$, where η is the norm of the sample mean vector defined in [37] as:

$$\eta = \sqrt{\left(\frac{1}{N} \sum_{i=1}^N \cos \theta_i\right)^2 + \left(\frac{1}{N} \sum_{i=1}^N \sin \theta_i\right)^2} \quad (7)$$

Being A a non invertible function, we rely on the well-known and qualified method proposed in [12], which approximates $A^{-1}(\eta)$ by:

$$\hat{\tau} = \tilde{A}^{-1}(\eta) = \begin{cases} 2\eta + \eta^3 + \frac{5\eta^5}{6} & \eta < 0.53 \\ -0.4 + 1.39\eta + \frac{0.43}{1-\eta} & 0.53 \leq \eta < 0.85 \\ \frac{1}{\eta^3 - 4\eta^2 + 3\eta} & \eta \geq 0.85 \end{cases} \quad (8)$$

Once an estimate of the **VM pdf** is obtained, we need to compare it with those computed on synthetic data of known **ids**. To this aim, a closed-form of the KL between two **VM pdfs** of parameters (ν_1, τ_1) and (ν_2, τ_2) is defined in [40] as:

$$\begin{aligned}\overline{KL}_{\nu, \tau} &= \mathcal{KL}(q(\cdot; \nu_1, \tau_1), q(\cdot; \nu_2, \tau_2)) = \int_{-\pi}^{\pi} q(\theta; \nu_1, \tau_1) \log \left(\frac{q(\theta; \nu_1, \tau_1)}{q(\theta; \nu_2, \tau_2)} \right) d\theta \\ &= \log \frac{I_0(\tau_2)}{I_0(\tau_1)} + \frac{I_1(\tau_1) - I_1(-\tau_1)}{2I_0(\tau_1)} (\tau_1 - \tau_2 \cos(\nu_2 - \nu_1))\end{aligned}\quad (9)$$

3.3 Combining Angle and Norm Concentration

In the previous sections we described the base theory laying foundations for an **id** estimator exploiting the information conveyed by the concentration of norms and angles. To provide a unique technique that combines these information, we should compare the joint empirical **pdf** $\hat{h}(r, \theta)$ related to the real dataset with the D theoretical **pdfs**, which will be referred to as $h_d(r, \theta)$ (where $d \in \{1..D\}$). Summarizing, the **id** estimate we want to compute is:

$$\hat{d} = \arg \min_{1 \leq d \leq D} \int_{-\pi}^{\pi} \int_0^1 h_d(r, \theta) \log \left(\frac{h_d(r, \theta)}{\hat{h}(r, \theta)} \right) dr d\theta$$

Note that neither the theoretical h_d is easily derivable, nor the joint **pdf** \hat{h} can be precisely estimated. Luckily, the norm distribution $g(r; k, d)$ and the angle distribution $q(\theta; \nu, \tau)$ are independent when the data are uniformly drawn from a spherical distribution [29]; therefore the joint **pdf** factorizes in the product of the two marginals, i.e. $h_d(r, \theta) = g(r; k, d)q(\theta; \nu, \tau)$, so that the KL divergence between $h_d(r, \theta)$ and $\hat{h}(r, \theta)$ becomes:

$$\overline{KL}_{d, \nu, \tau} = \mathcal{KL}(h_d(r, \theta), \hat{h}(r, \theta)) = \overline{KL}_d + \overline{KL}_{\nu, \tau} \quad (10)$$

This fact allows to split the joint $\overline{KL}_{d, \nu, \tau}$ in the sum of the two closed-form divergences reported in Equations (3, 9); it follows that the **id** estimator exploited in our algorithm becomes: $\hat{d} = \arg \min_{1 \leq d \leq D} \overline{KL}_{d, \nu, \tau}$.

4 The Algorithm

In this section we show how the theoretical results presented in Section 3 can be exploited to estimate the **id** of a given dataset combining the information expressed by the angles and by the minimum neighbor distances.

More precisely, we consider a manifold $\mathcal{M} \equiv \mathbb{R}^d$ embedded in a higher dimensional space \mathbb{R}^D through a locally isometric nonlinear smooth map $\phi : \mathcal{M} \rightarrow \mathbb{R}^D$, and a sample set $\mathbf{X}_N = \{\mathbf{x}_i\}_{i=1}^N = \{\phi(\mathbf{z}_i)\}_{i=1}^N \subset \mathbb{R}^D$, where \mathbf{z}_i are independent identically distributed points drawn from \mathcal{M} according to a non-uniform smooth pdf $f : \mathcal{M} \rightarrow \mathbb{R}^+$.

To estimate the **id** of \mathcal{M} , for each point $\mathbf{x}_i \in \mathbf{X}_N$ we find the set of $k+1$ ($1 \leq k \leq N-2$) nearest neighbors $\bar{\mathbf{X}}_{k+1} = \bar{\mathbf{X}}_{k+1}(\mathbf{x}_i) = \{\mathbf{x}_j\}_{j=1}^{k+1} \subset \mathbf{X}_N$. Calling $\hat{\mathbf{x}} = \hat{\mathbf{x}}_{k+1}(\mathbf{x}_i) \in \bar{\mathbf{X}}_{k+1}$ the farthest neighbor of \mathbf{x}_i , we calculate the distance between \mathbf{x}_i and its nearest neighbor in $\bar{\mathbf{X}}_{k+1}$, and we normalize it by means of the distance between \mathbf{x}_i and $\hat{\mathbf{x}}$. More precisely:

$$\rho(\mathbf{x}_i) = \min_{\mathbf{x}_j \in \bar{\mathbf{X}}_{k+1}} \frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\|\mathbf{x}_i - \hat{\mathbf{x}}\|} \quad (11)$$

This equation is used to compute a vector of normalized distances $\hat{\mathbf{r}} = \{\hat{r}_i\}_{i=1}^N = \{\rho(\mathbf{x}_i)\}_{i=1}^N$. By employing Equation (7) in [28], we compute the ML estimation by numerically solving the optimization problem $\hat{d}_{ML} = \arg \max_{1 \leq d \leq D} ll(d)$, where:

$$ll(d) = N \log kd + (d-1) \sum_{\mathbf{x}_i \in \mathbf{X}_N} \log \rho(\mathbf{x}_i) + (k-1) \sum_{\mathbf{x}_i \in \mathbf{X}_N} \log(1 - \rho^d(\mathbf{x}_i))$$

Similarly, for each point $\mathbf{x}_i \in \mathbf{X}_N$ we find its k nearest neighbors $\bar{\mathbf{X}}_k$ and we center them by means of a translation to obtain $\hat{\mathbf{X}}_k = \{\mathbf{x}_j - \mathbf{x}_i : \forall \mathbf{x}_j \in \bar{\mathbf{X}}_k\}$; next, we calculate $\binom{k}{2}$ angles of all the possible pairs of vectors in $\hat{\mathbf{X}}_k$, as follows:

$$\theta(\mathbf{x}_z, \mathbf{x}_j) = \arccos \frac{\mathbf{x}_z \cdot \mathbf{x}_j}{\|\mathbf{x}_z\| \|\mathbf{x}_j\|} \quad (12)$$

For each neighborhood we compute a vector $\hat{\boldsymbol{\theta}} = \{\theta(\mathbf{x}_z, \mathbf{x}_j)\}_{1 \leq i < j \leq k}$ by means of Equation (12). Since $\hat{\boldsymbol{\theta}}$ follows a VM pdf of parameters ν and τ (see Section 3.2), we estimate their values by employing the ML approach described in Equations (6, 8) for each set of neighbors, thus obtaining the vectors $\hat{\boldsymbol{\nu}} = \{\hat{\nu}_i\}_{i=1}^N$ and $\hat{\boldsymbol{\tau}} = \{\hat{\tau}_i\}_{i=1}^N$, and their means $\hat{\mu}_\nu = N^{-1} \sum_{i=1}^N \hat{\nu}_i$ and $\hat{\mu}_\tau = N^{-1} \sum_{i=1}^N \hat{\tau}_i$.

Moreover, for each dimensionality $d \in \{1..D\}$ we uniformly draw a set of N points $\mathbf{Y}_{Nd} = \{\mathbf{y}_i\}_{i=1}^N$ from the unit d -dimensional hypersphere², and we similarly compute a vector of normalized distances $\check{\mathbf{r}}_d = \{\check{r}_{id}\}_{i=1}^N = \{\rho(\mathbf{y}_i)\}_{i=1}^N$ and its ML estimation $\check{d}_{d,ML}$. Next, we calculate the vectors of the VM distribution parameters $\check{\boldsymbol{\nu}}_d = \{\check{\nu}_i\}_{i=1}^N$ and $\check{\boldsymbol{\tau}}_d = \{\check{\tau}_i\}_{i=1}^N$ together with their means $\check{\mu}_\nu^d$ and $\check{\mu}_\tau^d$.

Finally, we compose Equations (3, 9) as reported in Equation (10), thus obtaining the following **id** estimate:

$$\hat{d} = \arg \min_{d \in \{1..D\}} \mathcal{KL}(g(\cdot; k, \hat{d}_{ML}), g(\cdot; k, \check{d}_{d,ML})) + \mathcal{KL}(q(\cdot; \hat{\mu}_\nu, \hat{\mu}_\tau), q(\cdot; \check{\mu}_\nu^d, \check{\mu}_\tau^d)) \quad (13)$$

We call this **id** estimator **DANCo** (Dimensionality from Angle and Norm Concentration). Its time complexity is $O(D^2 N \log N)$ and it is dominated by the time complexity of the kNN algorithm ($O(DN \log N)$).

Considering Theorem 4 in [9], which ensures that geodetic distances in the infinitesimal ball converge to Euclidean distances with probability 1, and the results in Theorems 1, 2, Equation (13) represents a consistent estimator for the **id** of the manifold \mathcal{M} .

² Notice that a d -dimensional vector randomly sampled from a d -dimensional hypersphere according to the uniform pdf can be generated by drawing a point $\bar{\mathbf{y}}$ from a standard normal distribution $\mathcal{N}(\mathbf{0}_d, 1)$ and by scaling its norm.

5 Algorithm Evaluation

In this section we describe the datasets employed in our experiments (see Section 5.1), we summarize the adopted experimental settings (see Section 5.2), and we report the results achieved by the proposed algorithm, comparing them to those obtained by state of the art *id* estimators (see Section 5.3).

5.1 Dataset Description

To evaluate our algorithm, we have performed experiments on the 17 synthetic and 5 real datasets reported in Table 1. In details, to generate 15 synthetic datasets we have employed the tool proposed in [19], extending it to produce the datasets \mathcal{M}_{13} and \mathcal{M}_{14} by drawing points from nonlinearly embedded manifolds having high *id*. More precisely, to generate \mathcal{M}_{13} we have proceeded as follows: starting from 2500 points $\{\mathbf{x}_i\}_{i=1}^{2500}$ uniformly drawn in $[0, 1]^{18}$, we multiplied each \mathbf{x}_i first by $\sin(\cos(2\pi\mathbf{x}_i))$, then by $\cos(\sin(2\pi\mathbf{x}_i))$, obtaining points in $[0, 1]^{36}$ after a concatenation of the above coordinates. The dataset \mathcal{M}_{13} , containing 2500 points in $[0, 1]^{72}$, was finally obtained by duplicating each point's coordinate; this dataset, whose *id* is 18, is composed by points drawn from a manifold nonlinearly embedded in \mathbb{R}^{72} . The dataset \mathcal{M}_{14} was similarly generated starting from the same number of uniformly sampled points in \mathbb{R}^{24} .

The real datasets employed are: the ISOMAP face database [36], the MNIST database [25], the **Santa Fe** [32] dataset, the **Isolet** dataset [13], and the DSVIC1 time series [4].

The ISOMAP face database consists in 698 gray-level images of size 64×64 depicting the face of a sculpture. This dataset has three degrees of freedom: two for the pose and one for the lighting direction.

The MNIST database consists in 70000 gray-level images of size 28×28 of hand-written digits; in our tests we used the 6742 training points representing the digit 1. The *id* of this database is not actually known; we therefore rely on the estimations proposed in [19,9] for the different digits, and in particular on the range $\{8..11\}$ for the digit 1.

The version *D2* of the **Santa Fe** dataset is a synthetic time series of 50000 one-dimensional points; it was generated by a simulation of particle motion, and it has nine degrees of freedom. In order to estimate the attractor dimension of this time series, we used the method of delays described in [31], which generates *D*-dimensional vectors by collecting *D* values from the original dataset; by choosing $D = 50$ we obtained a dataset containing 1000 points in \mathbb{R}^{50} .

The **Isolet** dataset has been generated as follows: 150 subjects spoke the name of each letter of the alphabet twice, thus producing 52 training examples from each speaker. The latter are grouped into sets of 30 speakers each, and are referred to as *isolet1*, *isolet2*, *isolet3*, *isolet4*, and *isolet5*, for a total of 7797 samples. The *id* of this dataset is not actually known, but a study reported in [24] has proposed that the correct estimation could be in the range $\{16..22\}$.

Table 1. Brief description of the 17 synthetic and 5 real datasets employed in our experiments, where d is the `id` and D is the embedding space dimension.

Dataset	Name	d	D	Description
Synthetic	\mathcal{M}_1	10	11	Uniformly sampled sphere linearly embedded.
	\mathcal{M}_2	3	5	Affine space.
	\mathcal{M}_3	4	6	Concentrated figure, confusable with a $3d$ one.
	\mathcal{M}_4	4	8	Nonlinear manifold.
	\mathcal{M}_5	2	3	2-d Helix
	\mathcal{M}_6	6	36	Nonlinear manifold.
	\mathcal{M}_7	2	3	Swiss-Roll.
	\mathcal{M}_8	20	20	Affine space.
	\mathcal{M}_{9a}	10	11	Uniformly sampled hypercube.
	\mathcal{M}_{9b}	17	18	Uniformly sampled hypercube.
	\mathcal{M}_{9c}	24	25	Uniformly sampled hypercube.
	\mathcal{M}_{9d}	70	71	Uniformly sampled hypercube.
	\mathcal{M}_{10}	2	3	Möebius band 10-times twisted.
	\mathcal{M}_{11}	20	20	Isotropic multivariate Gaussian.
	\mathcal{M}_{12}	1	13	Curve.
	\mathcal{M}_{13}	18	72	Nonlinear manifold.
	\mathcal{M}_{14}	24	96	Nonlinear manifold.
Real	$\mathcal{M}_{\text{Faces}}$	3	4096	ISOMAP face dataset.
	$\mathcal{M}_{\text{MNIST1}}$	8 – 11	784	MNIST database (digit 1).
	$\mathcal{M}_{\text{SantaFe}}$	9	50	Santa Fe dataset (version $D2$).
	$\mathcal{M}_{\text{Isolet}}$	16 – 22	617	Spoken letter of the alphabet
	$\mathcal{M}_{\text{DSVC1}}$	2.26	20	Real time series of a Chua’s circuit.

The DSVC1 is a real data time series composed of 5000 samples and measured from a hardware realization of the Chua’s circuit [6]. We used the method of delays choosing $D = 20$, and we obtained a dataset containing 250 points in \mathbb{R}^{20} ; the `id` of this dataset is ~ 2.26 as reported in [4].

5.2 Experimental Setting

To objectively assess our method, we compared it with well-known `id` estimators such as: SPPCA, kNNG, CD, MLE, Hein, BPCA, MiND_{KL}, and IDEA. For kNNG, MLE, Hein, BPCA, MiND_{KL}, and IDEA we used the authors’ implementation³, while for the other algorithms we employed the version provided by the dimensionality reduction toolbox⁴.

To generate the synthetic datasets we adopted the modified generator described in Section 5.1 creating 20 instances of each dataset reported in Table 1, each of which is composed by 2500 randomly sampled points.

³ <http://www.eecs.umich.edu/~hero/IntrinsicDim/>,
<http://www.stat.lsa.umich.edu/~elevina/mledim.m>,
<http://www.ml.uni-saarland.de/code.shtml>,
<http://research.microsoft.com/en-us/um/cambridge/projects/infernet/blogs/bayesianpca.aspx>
<http://security.dico.unimi.it/~fox721/>

⁴ <http://cseweb.ucsd.edu/~lvdmaaten/dr/download.php>

Table 2. Parameter settings for the different estimators: k represents the number of neighbors, γ is the edge weighting factor for **kNNG**, M is the number of Least Square (LS) runs, N is the number of resampling trials per LS iteration, α and π represent the parameters (shape and rate) of the Gamma prior distributions describing the hyperparameters and the observation noise model of **BPCA**, μ contains the mean and the precision of the Gaussian prior distribution describing the bias inserted in the inference of **BPCA**.

Dataset	Method	Parameters
Synthetic	SPPCA	<i>None</i>
	CD	<i>None</i>
	MLE	$k_1 = 6 \ k_2 = 20$
	kNNG ₁	$k_1 = 6, k_2 = 20, \gamma = 1, M = 1, N = 10$
	kNNG ₂	$k_1 = 6, k_2 = 20, \gamma = 1, M = 10, N = 1$
	BPCA	$iters = 500, \alpha = (2.0, 2.0) \ \pi = (2.0, 2.0) \ \mu = (0.0, 0.01)$
	MiND _{KL}	$k = 10$
	IDEA	$k = 10$
	DANCo	$k = 10$
Real	SPPCA	<i>None</i>
	CD	<i>None</i>
	MLE	$k_1 = 3 \ k_2 = 8$
	kNNG ₁	$k_1 = 3, k_2 = 8, \gamma = 1, M = 1, N = 10$
	kNNG ₂	$k_1 = 3, k_2 = 8, \gamma = 1, M = 10, N = 1$
	BPCA	$iters = 2000, \alpha = (2.0, 2.0) \ \pi = (2.0, 2.0) \ \mu = (0.0, 0.01)$
	MiND _{KL}	$k = 5$
	IDEA	$k = 5$
	DANCo	$k = 5$

To obtain an unbiased estimation, for each technique we averaged the results achieved on the 20 instances. To execute multiple tests also on $\mathcal{M}_{\text{MNIST1}}$ and $\mathcal{M}_{\text{Isolet}}$ we extracted 5 random subsets containing 2500 points each, and we averaged the achieved results.

In Table 2 the configuration parameters employed in our tests are summarized. To relax the dependency of the **kNNG** algorithm from the selection of the value of its parameter k , we performed multiple runs with $k_1 \leq k \leq k_2$ and we averaged the achieved results (see Table 2).

5.3 Experimental Results

This section reports the results achieved on both synthetic and real datasets. In particular, Table 3 summarizes the results obtained on the synthetic datasets. It is possible to note that the best performing algorithm is **DANCo**. Indeed, this estimator can correctly deal with linear and nonlinear manifolds embedded in low and high dimensional spaces. In particular, it is the only method that achieves a good estimation for the three datasets \mathcal{M}_{9d} , \mathcal{M}_{13} , and \mathcal{M}_{14} .

Instead, geometrical approaches, such as **kNNG**, **CD**, **MLE**, and **Hein**, obtain good estimates only for low **id** manifolds, failing to deal with high **id** data.

Table 3. Results achieved on the synthetic datasets. The best approximations are highlighted in boldface.

Dataset	d	SPPCA	BPCA	kNNG ₁	kNNG ₂	CD	MLE	Hein	MiND _{KL}	IDEA	DANCo
\mathcal{M}_{12}	1	3.00	5.70	0.97	1.07	1.14	1.00	1.00	1.00	1.02	1.00
\mathcal{M}_5	2	3.00	2.00	1.96	2.06	1.98	1.97	2.00	2.00	2.00	2.00
\mathcal{M}_7	2	3.00	2.00	1.97	2.09	1.93	1.96	2.00	2.00	2.07	2.00
\mathcal{M}_{10}	2	3.00	1.55	1.95	2.03	2.19	2.21	2.00	2.00	1.98	2.00
\mathcal{M}_2	3	3.00	3.00	2.95	3.03	2.88	2.88	3.00	3.00	3.03	3.00
\mathcal{M}_3	4	4.00	4.00	3.75	3.82	3.23	3.83	4.00	4.00	4.01	4.00
\mathcal{M}_4	4	8.00	4.25	4.05	4.76	3.88	3.95	4.00	4.15	3.93	4.00
\mathcal{M}_6	6	12.00	12.00	6.46	11.24	5.91	6.39	5.95	6.50	6.33	6.90
\mathcal{M}_1	10	11.00	5.45	9.16	9.89	9.12	9.10	9.45	10.30	10.41	10.00
\mathcal{M}_{9a}	10	10.00	5.20	8.62	10.21	8.09	8.26	8.90	9.85	9.93	9.50
\mathcal{M}_{9b}	17	17.00	9.46	13.69	15.38	12.30	12.87	13.85	16.25	16.07	16.47
\mathcal{M}_{13}	18	36.00	36.00	17.58	5.01	11.60	15.95	14.00	18.60	17.30	18.20
\mathcal{M}_8	20	20.00	13.55	15.25	10.59	13.75	14.64	15.50	19.15	18.51	19.54
\mathcal{M}_{11}	20	20.00	13.70	16.40	24.89	11.26	15.82	15.00	19.35	21.20	19.90
\mathcal{M}_{9c}	24	24.00	13.3	17.67	21.42	15.58	16.96	17.95	22.55	23.93	23.85
\mathcal{M}_{14}	24	48.00	48.00	19.66	22.80	14.03	19.83	17.00	25.30	22.90	25.00
\mathcal{M}_{9d}	70	71.00	71.00	39.67	40.31	31.4	36.49	38.69	65.30	46.7	70.42
MPE		44.79	61.55	11.72	20.14	20.79	13.78	12.04	2.94	4.75	1.90

Moreover, the projection techniques, such as BPCA and SPPCA, are able to correctly deal only with linear embedded manifolds. These considerations confirm that the geometric methods are affected by an underestimation bias as noticed in [28,34] and that all the projection methods cannot provide reliable id estimates [26].

Furthermore, DANCo outperforms also IDEA and MiND_{KL} that have been developed to deal with datasets having a sufficiently high id (that is $\text{id} \geq 10$) and being drawn from manifolds nonlinearly embedded in higher dimensional spaces.

In the last row of Table 3 the Mean Percentage Error (MPE) indicator, proposed in [28] in order to evaluate the overall performance of a given estimator, is reported. For each algorithm this value is computed as the mean of the percentage errors obtained on each dataset, i.e. $\text{MPE} = \frac{100}{\#\mathcal{M}} \sum_{\mathcal{M}} \frac{|\hat{d}_{\mathcal{M}} - d_{\mathcal{M}}|}{d_{\mathcal{M}}}$, where $d_{\mathcal{M}}$ is the real id, $\hat{d}_{\mathcal{M}}$ is the estimated one, and $\#\mathcal{M}$ is the number of tested manifolds. Considering this indicator, DANCo ranks as the best performing estimator.

In Table 4 the results obtained on real datasets are summarized. Being the real data generally affected by the presence of noise, the quality of the estimates computed by the projection methods is strongly reduced, as confirmed by the poor results obtained by BPCA and SPPCA. The geometric approaches we tested are less affected by noise, but they are not able to correctly deal with the high dimensionality of the $\mathcal{M}_{\text{Isolet}}$ dataset.

As can be seen, DANCo is the best performing estimator, strongly overcoming also the results obtained by those techniques, such as IDEA and MiND_{KL}, that exploit a correction approach. These results, together with the best average

Table 4. Results achieved on the real datasets by the employed approaches. The best approximations are highlighted in boldface.

Dataset	d	SPPCA	BPCA	kNNG ₁	kNNG ₂	CD	MLE	Hein	MinD _{KL}	IDEA	DANCo
$\mathcal{M}_{\text{DSVC1}}$	2-26	4.00	6.00	1.77	1.86	1.92	2.03	3.00	2.50	2.14	2.26
$\mathcal{M}_{\text{Faces}}$	3	5.00	4.00	3.60	4.32	3.37	4.05	3.00	3.90	3.73	4.00
$\mathcal{M}_{\text{Santa Fe}}$	9	19.00	18.00	7.28	7.43	4.39	7.16	6.00	7.60	7.26	8.19
$\mathcal{M}_{\text{MNIST1}}$	8-11	9.00	11.00	10.37	9.58	6.96	10.29	8.00	11.00	11.06	9.98
$\mathcal{M}_{\text{Isolet}}$	16-22	45.00	19.00	6.50	8.32	3.65	15.78	3.00	20.00	18.77	19.00
MPE		79.37	62.92	27.14	27.24	37.22	18.17	33.21	15.44	13.32	9.47

estimation precision achieved by our technique in terms of MPE^5 , confirm that DANCo is a promising and valuable tool for id estimation.

Finally, to test the robustness of our algorithms w.r.t. the choice of the parameter k , we employed DANCo to reproduce the experiments proposed for the MLE algorithm in Figure 1 (a) of [26] and in Figure 2 of [28], and we averaged the curves obtained in 10 runs. In these tests the adopted datasets are composed by points drawn from the standard Gaussian pdf in \mathbb{R}^5 . We repeated the test for datasets with cardinalities $N \in \{200, 500, 1000, 2000\}$ varying the parameter k in the range $\{5..100\}$. For all the combinations of the dataset cardinalities and the k parameter values, DANCo obtained id estimates always equal to 5, confirming its strong robustness.

6 Conclusions and Future Works

In this paper we proposed a novel consistent estimator, called DANCo, that combines the effects of concentration of angles and norms to estimate the id of a given dataset. The proposed method compares the joint pdf estimated on the dataset, related to angles and norms respectively, with those computed on synthetic datasets of known id; to this aim, a closed-form expression for the Kullback-Leibler divergence of their distributions is employed.

We tested our algorithm on both synthetic and real datasets comparing its results with those obtained by employing well-known id estimators. The overall results show that DANCo is a really promising and valuable technique for id estimation. Indeed, it provides the most accurate results, computing either the best id estimates or values that are strongly comparable to the best ones. Moreover, this algorithm has shown to be really robust in terms of its capability to: i) deal with both high and low id, ii) manage both linearly and nonlinearly embedded manifolds, and iii) outperform all the other estimators on noisy real datasets.

Future works will be devoted to identify a bound for the finite sample error, to further formally evaluate the effectiveness of the proposed approach.

⁵ Where the true value of the id is not known, we considered the mean value of the range as $d_{\mathcal{M}}$.

References

1. C.M. Bishop. Bayesian PCA. *Proc. of NIPS*, 11:382–388, 1998.
2. E. Breitenberger. Analogues of the normal distribution on the circle and the sphere. *Biometrika*, 50, 1963.
3. F. Camastra. Data dimensionality estimation methods: A survey. *Pattern Recognition*, 36(12):2945–2954, 2003.
4. F. Camastra and M. Filippone. A comparative evaluation of nonlinear dynamics methods for time series prediction. *Neural Computing and Applications*, 18(8):1021–1029, November 2009.
5. F. Camastra and A. Vinciarelli. Estimating the intrinsic dimension of data with a fractal-based method. *IEEE Trans. on PAMI*, 24:1404–1407, 2002.
6. L. Chua, M. Komuro, and T. Matsumoto. The double scroll. *IEEE Trans. on Circuits and Systems*, 32:797–818, 1985.
7. J.A. Costa and A.O. Hero. Geodesic entropic graphs for dimension and entropy estimation in manifold learning. *IEEE Trans. on Signal Processing*, 52(8):2210–2221, 2004.
8. J.A. Costa and A.O. Hero. Learning intrinsic dimension and entropy of high-dimensional shape spaces. In *Proc. of EUSIPCO*, 2004.
9. J.A. Costa and A.O. Hero. Learning intrinsic dimension and entropy of shapes. In *Stat. and anal. of shapes*. Birkhauser, 2005.
10. J.P. Eckmann and D. Ruelle. Fundamental limitations for estimating dimensions and Lyapunov exponents in dynamical systems. *Physica D: Nonlinear Phenomena*, 56(2-3):185–187, 1992.
11. A.M. Farahmand, C. Szepesvari, and J.Y. Audibert. Manifold-adaptive dimension estimation. *Proc. of ICML*, 2007.
12. N. I. Fisher. *Statistical Analysis of Circular Data*. Cambridge University Press, January 1996.
13. A. Frank and A. Asuncion. UCI machine learning repository, 2010.
14. J.H. Friedman, T. Hastie, and R. Tibshirani. *The Elements of Statistical Learning - Data Mining, Inference and Prediction*. Springer, Berlin, 2009.
15. K. Fukunaga. An algorithm for finding intrinsic dimensionality of data. *IEEE Trans. on Computers*, 20:176–183, 1971.
16. K. Fukunaga. *Intrinsic Dimensionality Extraction*. Classification, Pattern Recognition and Reduction of Dimensionality. P.R. Krishnaiah and L.N. Kanal, Amsterdam: North Holland, 1982.
17. P. Grassberger and I. Procaccia. Measuring the strangeness of strange attractors. *Physica D: Nonlinear Phenomena*, 9:189–208, 1983.
18. Y. Guan and J. G. Dy. Sparse probabilistic principal component analysis. *J. of Machine Learning Research - Proc. Track*, 5:185–192, 2009.
19. M. Hein. Intrinsic dimensionality estimation of submanifolds in euclidean space. In *Proc. of ICML*, pages 289–296, 2005.
20. G. W. Hill. New approximations to the von Mises distribution. *Biometrika*, 63(3):673–676, 1976.
21. I.T. Jolliffe. *Adaptive Control Processes: A Guided Tour*. Princeton University Press, 1961.
22. I.T. Jolliffe. *Principal Component Analysis*. Springer Series in Statistics. Springer-Verlag, New York, NY, 1986.
23. M. Kirby. *Geometric Data Analysis: an Empirical Approach to Dimensionality Reduction and the Study of Patterns*. John Wiley and Sons, 1998.

24. I. Kivimäki, K. Lagus, I. Nieminen, J. Väyrynen, and T. Honkela. Using correlation dimension for analysing text data. In *Proc. of the ICANN*, pages 368–373, Berlin, Heidelberg, 2010. Springer-Verlag.
25. Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proc. of IEEE*, 86:2278–2324, 1998.
26. E. Levina and P.J. Bickel. Maximum likelihood estimation of intrinsic dimension. *Proc. of NIPS 17*, 1:777–784, 2005.
27. J. Li and D. Tao. Simple exponential family PCA. *Proc. of AISTATS*, pages 453–460, 2010.
28. G. Lombardi, A. Rozza, C. Ceruti, E. Casiraghi, and P. Campadelli. Minimum neighbor distance estimators of intrinsic dimension. *Proc. of ECML-PKDD*, 6912:374–389, 2011.
29. R. D. Lord. The use of the Hankel transform in statistics I. general theory and examples. *Biometrika*, 41(1/2):44–55, 1954.
30. K. V. Mardia. *Statistics of Directional Data*. Academic Press, 1972.
31. E. Ott. *Chaos in Dynamical Systems*. Cambridge University Press, Cambridge, 1993.
32. F.J. Pineda and J.C. Sommerer. Estimating generalized dimensions and choosing time delays: A fast algorithm. *Time Series Prediction. Forecasting the Future and Understanding the Past*, pages 367–385, 1994.
33. A. Rozza, G. Lombardi, C. Ceruti, E. Casiraghi, and P. Campadelli. Novel high intrinsic dimensionality estimators. *Machine Learning J.*, May 2012.
34. A. Rozza, G. Lombardi, M. Rosa, E. Casiraghi, and P. Campadelli. IDEA: Intrinsic dimension estimation algorithm. *Proc. ICIAP*, 6978:433–442, 2011.
35. A. Sodergren. On the distribution of angles between the N shortest vectors in a random lattice. *J. London Math. Soc.*, 84(3):749–764, 2011.
36. J. Tenenbaum, V. Silva, and J. Langford. A global geometric framework for non-linear dimensionality reduction. *Science*, 290:2319–2323, 2000.
37. G. J. G. Upton. Approximate confidence intervals for the mean direction of a von Mises distribution. *Biometrika*, 73(2):525–527, 1986.
38. V. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, 1998.
39. P.J. Verveer and R.P.W. Duin. An evaluation of intrinsic dimensionality estimators. *IEEE Trans. on PAMI*, 17:81–86, 1995.
40. A.P.N. Vo, S. Orintara, and T.T. Nguyen. Statistical image modeling using von Mises distribution in the complex directional wavelet domain. In *Proc. of ISCAS 2008*, pages 2885–2888, 2008.
41. Q. Wang, S.R. Kulkarni, and S. Verd. A nearest-neighbor approach to estimating divergence between continuous random vector. In *Proc. ISIT*, pages 242–246, 2006.